A Comparison of Unconstraining Methods to Improve Revenue Management Systems

Carrie Crystal[a]
Mark Ferguson[a]*
Jon Higbie[b]
Rohit Kapoor[a]

[a]The College of Management
Georgia Institute of Technology
800 West Peachtree Street
Atlanta, GA 30332-0520

[b]Manugistics, Inc
2839 Paces Ferry Road, SE, Suite 1000
Atlanta, GA 30339-5770

*Corresponding Author: Tel: (404) 894-4330, Fax: (404) 894-6030, Mark.Ferguson@mgt.gatech.edu

A Comparison of Unconstraining Methods to Improve Revenue Management Systems

Abstract

A successful revenue management system requires accurate demand forecasts for each customer segment. These forecasts are used to set booking limits for lower value customers to ensure an adequate supply for higher value customers. The very use of booking limits, however, constrains the historical demand data needed for an accurate forecast. Ignoring this interaction leads to substantial penalties in a firm's potential revenues. We review existing unconstraining methods and propose a new method that includes some attractive properties not found in the existing methods. We evaluate several of the common methods used to unconstrain historical demand data against our proposed method by testing them on intentionally constrained simulated data. Results show our proposed method along with the Expectation Maximization (EM) method perform the best. We also test the revenue impact of our proposed method, EM, and "no unconstraining" on actual booking data from a hotel/casino. We show that performance varies with the initial starting protection limits and a lack of unconstraining leads to significant revenue losses.

Keywords: Revenue Management, Truncated Demand, Forecasting, Unconstrained Demand

1. Introduction

Revenue Management has been credited with improving revenues 3%-7% in the airline, hotel, and car rental industries (Cross, 1997). One of the core concepts behind revenue management is the reservation of a portion of capacity for higher value customers at a later date. The amount of capacity to reserve is typically determined through the calculation of booking limits, which place restrictions on the amount of capacity made available to a lower value segment of customers so as to reserve capacity for a higher value segment that may arrive in the future. Most booking limit calculations depend on the deduction of a demand distribution for each customer value segment from past demand data that occurred under similar circumstances and operating environments. In practice, however, true demand data is difficult to obtain as many firms are unable to record all demand request that arrive after a booking limit has been exceeded and capacity for that customer segment has been restricted.

To overcome this problem, "unconstraining" methods are used to extrapolate the true demand distribution parameters from truncated demand data collected over previous selling opportunities. Once a firm sells out of capacity for a given segment, the sales data for that segment represents truncated demand (equal to the booking limit) instead of true demand. While there is no perfect way to unconstrain sales data, Weatherford and Polt (2002) claim that, in the airline industry, switching from one common industry method to a better method increases revenues 0.5 to 1.0 percent. Since most firms using revenue management have low marginal costs, maximizing revenues translates into maximizing operating profits. Hence, unconstraining methods significantly impact revenues, and in turn, profits, and deserve closer research attention.

Despite the significant impact that unconstraining has on the success of a revenue management application, this topic has received much less attention in the literature compared to the work on methods for setting and adjusting booking limits. This is surprising since the demand distribution parameter estimates represent a primary input to most booking limit techniques, fundamentally linking the value of the former with the quality of the latter. A firm facing constrained sales data faces three choices: 1) leave the data constrained, 2) directly observe and record latent demand, or 3) statistically unconstrain the data after the fact.

If historical sales data is left constrained, true demand is underestimated, creating a spiral-down effect on total revenue where the firm's expected revenue decreases monotonically over time (Cooper et al. 2005). Unfortunately, due in part to the absence of research and teaching in this area, this practice is common at firms using less sophisticated revenue management systems. In §5 we demonstrate how ignoring constrained data impacts revenue using actual booking data from a hotel/casino.

Direct observation involves the recording of latent (unsatisfied) demand. Many hotels for example, record both bookings (requests that are met) and turndowns (requests that are not met). Care must be taken however, as turndowns may be attributed to availability (denials) or price (regrets). The former is considered latent demand while the latter is not. Many hotel chains have invested in systems and training for their reservations agents in order to track turndowns, and relied on these direct observations to unconstrain their sales data. Unfortunately, there are many issues with using turndown data for unconstraining including multiple availability inquiries from the same customer, incorrect categorization of turndowns by reservations agents, and the fact that only small potions of customer request arrive through a channel controlled by the firm (Orkin 1998). The latter provides the largest hurdle for most industries.

Direct observations of demand are not an option for many industries because of their distribution channels. For example, traditionally, most airline bookings have been made via travel agents using global distribution systems like Sabre and Worldspan, and no turndown information is collected on these bookings. While airlines have recently been striving to increase their direct sales and improve their customer information, the percentage of total demand collected through these channels is still very small. On the other hand, hotels and casinos have historically taken the majority of their bookings through their own agents, either at the property itself or through a central reservations center. The advent of the Internet however has compromised the quality of their turndown data. While Internet direct sales is a growing channel for hotels, with some hotels taking up to 10% of their booking through this channel, most companies have yet to incorporate turndowns from their own web site into their total demand picture, and for good reason. Carroll and Siguaw (2003) point out that only 20% of hotel customers that check availability via the direct Internet channel actually book their rooms at that same site. Along with the growth in direct Internet sales, sales via third party web sites (such as Expedia and Travelocity) have grown at an even faster rate. Most third party web sites do not provide any turndown information. The net effect for hotels is an increasing proportion of bookings from channels for which no turndowns are collected, and as a result, hoteliers have increasing interest in alternative unconstraining methods.

Statistical unconstraining covers a spectrum of optimization and heuristic techniques that rely only on observed bookings and a state indicator (open/closed). The purpose of this paper is to compare some of the most common statistical unconstraining methods that have appeared in the literature and compare them to our proposed forecasting-based method. In addition, we apply the most accurate of these methods to real hotel booking data. Previous studies on unconstraining methods have only tested a subset of methods against simulated airline data.

Most traditional unconstraining methods follow a similar methodology. They label the demand observed over time for a given section of the booking period into constrained and unconstrained categories, then adjust the parameter value estimates for the demand distribution based on the percentage of data that was constrained. These unconstraining methods treat all truncated demand the same, a demand stream that is constrained twenty days before the end of the booking window is treated the same as one that is constrained one day before. This methodology ignores an important aspect of the revenue management environment: firms often know the time when demand is constrained. Our proposed method takes advantage of this information and uses it when calculating the demand distribution parameter estimates. In addition, our method offers two other advantages over many of the alternatives. It is based on a widely accepted statistical forecasting technique (double exponential smoothing) requiring minimal computations and it is non-parametric, requiring no a-priori assumptions about the shape of the booking curve or the distribution of the final total demand.

Figure 1 - Example booking pace curve

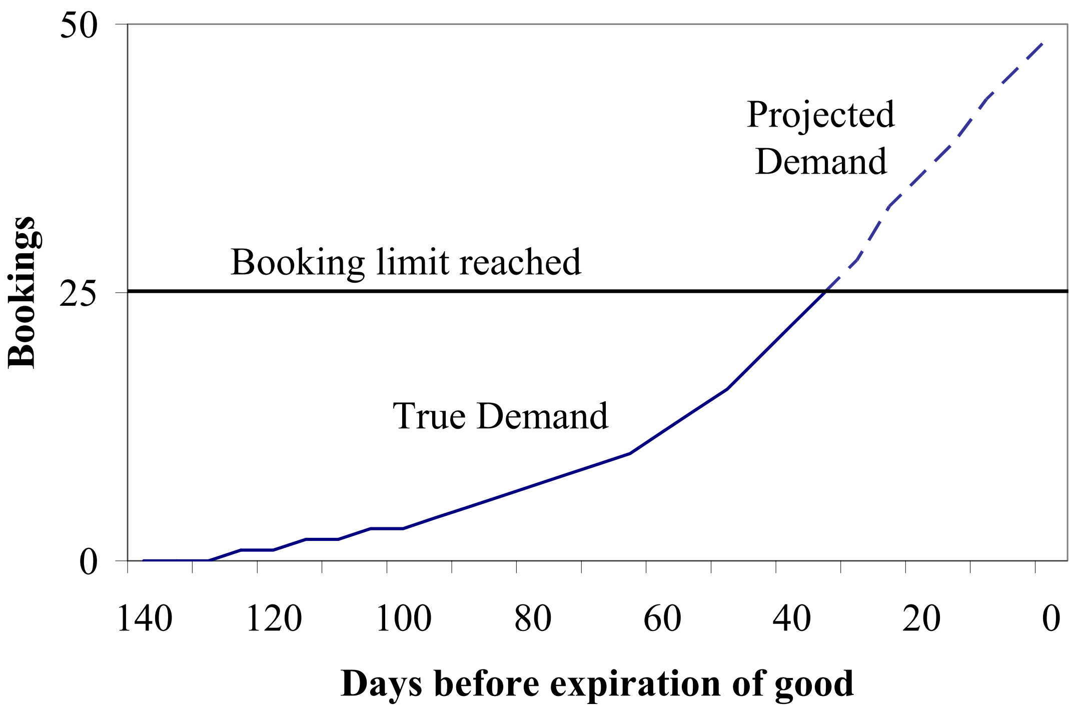

We illustrate the key concept behind our proposed method through the example booking curve shown in Figure 1. Most traditional unconstraining methods only use the fact that demand was truncated and the final demand observed was 25, the booking limit for this particular scenario. Our proposed method also uses the fact that demand was truncated 30 days before the plane departure or the guest arrival. Our method then uses double exponential smoothing to project the total demand that would have been observed in the absence of any booking limits (the dashed part of the booking pace curve). Through a simulation experiment, we find that our method outperforms the majority of the traditional statistical methods in estimating the demand distribution parameters of constrained data sets. Compared to the one method it does not always outperform, our method is simpler and works under conditions where the other method does not, such as when all historical data sets are constrained. Since there is no clear dominance by either method, we evaluate the impact on revenue that the two methods provide using actual booking and revenue data from a leading hotel/casino.

The rest of the paper is organized as follows. §2 reviews the literature, §3 defines the proposed method, §4 tests the method against other common methods used in practice, §5 tests the two best performing methods on real hotel/casino data and measures their impact on the hotel's revenue. Finally, §6 concludes the paper.

2. Literature Review

Weatherford and Bodily (1992) and McGill and van Ryzin (1999) provide general reviews of the broad range of literature in the revenue management field. As these studies show, the primary research focus has been weighted towards the development of overbooking and booking limit techniques with little focus on unconstraining (also called uncensoring) sales data. We concentrate here on reviewing the unconstraining research.

Reliability engineers, biomedical scientists and econometricians have used unconstraining procedures for many years to compensate for the early termination of experiments. This is similar to the way revenue managers "terminate" demand for a particular customer value segment through the use of booking limits. Relevant research in these fields include: (Cox, 1972; Kalbfleisch and Prentice, 1980; Lawless, 1982; Cox and Oakes, 1984; Schneider, 1986; Nelson, 1990; Liu and Makis, 1996). These methods heavily rely on the use of the hazard rate function to determine the probability distribution of lifetime data. To our knowledge, van Ryzin and McGill (2000) provide the only use of this type method in a revenue management framework when they utilize a method based on demand lifetables. Lawless (1982) explains the lifetable method of uncensoring data and we include it in our comparison.

For more traditional revenue management unconstraining methods, Weatherford and Polt (2002) and Zeni (2001) compare unconstraining methods using simulation and apply the best methods to an airline's reservation data to test the revenue impact of using an inferior method. Six unconstraining methods are tested: three different averaging methods, booking profile (BP), projection detruncation (PD), and expectation maximization (EM). The averaging methods are the simplest computationally and therefore are often used in practice. We compare our proposed method (DES) against the three best performing methods found in Weatherford and Polt (2002): a simple averaging method (referred to as Naïve 3 in Weatherford and Polt, abbreviated to AM in this paper), EM and PD. Both Weatherford and Polt (2002) and Zeni (2001) conclude that the EM method outperforms the others and increases revenues by 2-12 % in full capacity situations. We also find the EM method provides the most accurate parameter estimates although there are many conditions where our proposed method is more accurate.

Of the three best methods that Weatherford and Polt (2002) and Zeni (2001) use, only EM and PD are grounded in statistical theory. Dempster et al. (1977) prove the theory behind the EM method based on data from a univariate distribution. The EM method discussed by Dempster et al. (1977) is essentially the same as the tobit model used in econometrics (Maddala, 1983). McGill (1995) extends the EM method to a multivariate problem when demand for different classes (segments) of a good are correlated. The PD method closely resembles the EM method, but takes a conditional median in place of a conditional mean. Additionally, the PD method allows users to change a weighting constant to obtain more aggressive demand estimates. The tradeoffs include increased computations and the risk of no solution convergence (Weatherford and Polt, 2002).

Liu et al. (2002) examine unconstraining demand data through the lens of the hotel industry and argue that the EM method is unrealistic in application because of its computational intensity. The authors argue that parametric regression models take into account all relevant information and are computationally more feasible in real-world applications. They develop a parametric regression model which uses booking curve data, but requires knowledge of the shape of the demand distribution and other specifics of the demand constraints. This knowledge requirement restricts the general use of their model, as firms often do not know a priori the shape of the booking curve. Also, the authors do not provide comparisons between their proposed parametric method and the methods discussed in other papers. We do not include their method in our comparison because we do not assume a known, functional form for the booking curve. We do agree, however with their criticism about the computational intensity required of the EM method. Our proposed method is much easier to calculate but, unlike the parametric models, does not require knowledge about the shape of the booking curves.

To test the revenue impact of our new unconstraining method, we must set the protection levels effectively. To do so, we use the most common seat protection heuristic used in practice, EMSR-b (Belobaba 1989). McGill and van Ryzin (1999) give an explanation of the EMSR-b method along with a review of the booking limits problem in general. Talluri and van Ryzin (2004b) provide an excellent treatment of all aspects of revenue management systems.

3. Proposed Unconstraining Method

Our proposed method uses Double Exponential Smoothing (DES) or Holt's Method to forecast the constrained values of a given data set. DES uses two smoothing constants, one for smoothing the base component of the demand pattern and a second for smoothing the trend component. Armstrong

(2001) provides a good review of this method. We describe how it may be used to solve the unconstraining problem below.

Let t represent the time periods between I, the period that reservations are initially accepted, and B, the period where demand reaches the booking limit (time is counted backwards). That is

$$t \in [I, I-1, \ldots B], I \geq B.$$

After period B, demand continues to occur but is unobserved. If demand is not constrained then $B = 0$. Thus, "demand seen" equates to the cumulative demand observed from periods I to B and is always less than or equal to true demand. From our example given in Figure 1, I corresponds to period 140 and B corresponds to period 30, after which demand is unobserved. Now define:

A_t = Actual demand in period t

F_t = The exponentially smoothed forecast for period t

T_t = The exponentially smoothed trend for period t

FIT_t = The forecast including trend for period t

α = Base smoothing constant

β = Trend smoothing constant

The forecast for the upcoming period t is

$$FIT_t = F_t + T_t \qquad (1)$$

where

$$F_t = FIT_{t-1} + \alpha(A_{t-1} - FIT_{t-1}) \qquad (2)$$

$$T_t = T_{t-1} + \beta(F_t - FIT_{t-1}). \qquad (3)$$

The smoothing constants, α and β, are decision variables. For each constrained demand instance, we use a non-linear optimization routine to select the alpha and beta values that minimize the sum of the squares of the forecast error:

$$\min_{\alpha,\beta} \sum_{t=I}^{B} (A_t - FIT_t)^2 \tag{4}$$

For the initial values, F_I and T_I, we use the actual demand in period I as our estimate for the base component and the average trend over the available data set as our estimate for the initial slope component. Since the problem is not jointly convex in α and β, a non-linear search algorithm such as tabu search or simulated annealing is needed to find the global minimizers. We then use our forecasting model to project the total demand over the number of periods that the data set is constrained. We do so by using (1) to forecast demand over periods B-1 to 0. Referring back to Figure 1, we calculate the α and β parameters over the solid portion (periods 140 to 30) of the demand curve and project the trend component over the dotted portion (periods 29 to 0) of the demand curve. The cumulative demand over the observed and projected components of the booking curve is then used as a single point estimate of true cumulative demand for a particular selling occurrence (i.e. a given Thursday night stay for a given rate program at a hotel). Call this individual point estimate for the ith booking curve X_i. We repeat this procedure over each constrained booking curve in a given data set (i.e. all Thursday night stays for a given rate program at a hotel). Thus, if there are n historical booking curves in the data set, we end up with a set of point estimates $(X_1, X_2, \ldots X_n)$. The final demand distribution parameters (mean μ and variance σ^2) are then estimated using this set of point estimates by:

$$\bar{\mu} = \frac{\sum_{i=1}^{n} X_i}{n} \quad \text{and} \quad \bar{\sigma}^2 = \frac{\sum_{i=1}^{n} (X_i - \mu)^2}{n} \, . \tag{5}$$

The basic model of DES described above is a very general method for forecasting demand and, as presented, does not account for seasonality, intermittent demand, opening and closing of booking limits, and other specifics that might be relevant in application. However, DES can be easily adjusted to incorporate these specific characteristics (Armstrong, 2001).

4. Comparison of Unconstraining Methods

In this section, we present our methodology for comparing five of the most common statistical unconstraining methods and the ensuing results. To compare the performance of the different methods, we simulate booking curves representing true demand and then impose booking limits to create constrained data. We apply five different unconstraining methods to the constrained data sets and compare the estimated demand parameters against the true parameters. The unconstraining method that estimates the demand distribution parameters closest to the actual true parameters is judged the best method.

To compare the performances of the chosen methods, we first simulate booking curves and set booking limits to constrain the data. To test each unconstraining method against a broad range of demand scenarios, we simulate three data sets with 100 booking curves each and 140 days in each booking curve. The three data sets represent three common shapes of booking curves: convex, homogeneous and concave (Liu et al., 2002) as shown in Figure 2. The 100 booking curves represent 100 historical demand records (for each shape curve) that a hotel or airline may use to predict future demand. For example, a hotel may keep demand data from its last 100 Thursday night stays in order to estimate demand for future Thursday night stays. Since most hotels and airlines see the great majority of their reservations within 140 days before the day of arrival or departure, we simulate 140 days of daily demand arrivals for each booking curve.; resulting in 100 individual booking curves of 140 days each, or 14,000 individual data points. For each booking curve shape, we looked at the total demand seen for all 100 booking curves simultaneously, (some where the total demand was not constrained and others where total demand exceeded the booking limits), and used each unconstraining method to estimate the true demand distribution parameters.

Figure 2 - Concave, Homogeneous and Convex booking curves

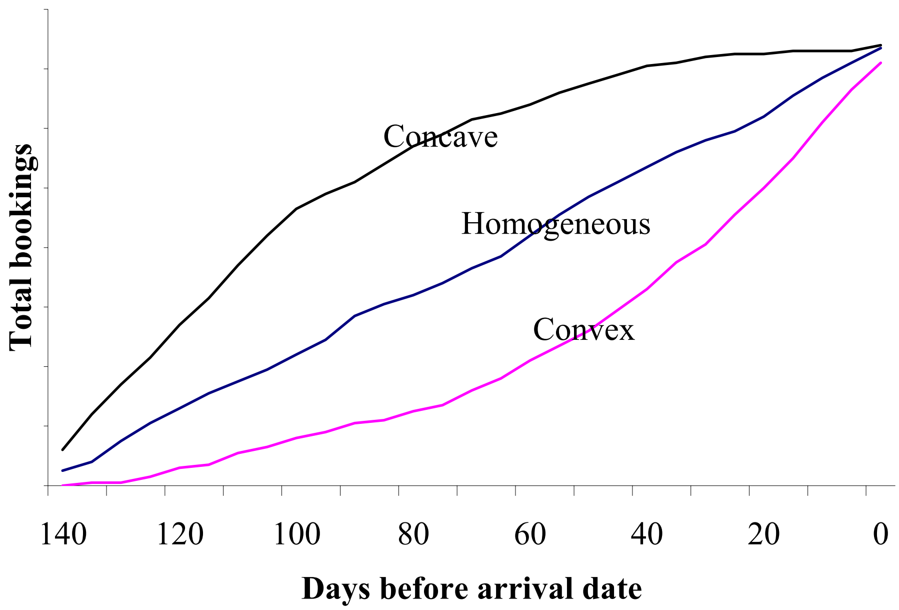

To construct the booking curve, we assume arrivals on a given day are randomly drawn from a Poisson distribution. This assumption is common in the literature and matches closely with actual data from the hotel and airline industries (Rothstein, 1974; Bitran and Mondschein, 1995; Bitran and Gilbert, 1996; Badinelli, 2000; Liu et al., 2002). For the homogeneous booking curve, we maintain a constant mean arrival rate over all 140 days. For the convex (concave) booking curves, we increment the mean arrival rate from low to high (high to low) respectively, so that the expected total demand over the 140 day period is the same for all three curves.

After we create the demand curves, we calculate booking limits. The minimum of true demand and the booking limit is the demand seen by the user. A simple example is shown in Table 1.

Table 1 - Example of true demand vs. demand observed

True Demand	100	110	91	95	103
Booking Limit	98	105	103	99	102
Demand Seen	98*	105*	91	95	102*

* indicates constrained demand, also called a closed segment

We use the fact that if daily demand arrivals are Poisson and the demands on different days are independent, then total demand is again Poisson. Since the mean of the Poisson-distributed total demand is sufficiently large, the distribution of total demand is approximately Normal. Thus, we calculate an expected average (μ) and standard deviation (σ) of the total demand and generate five sets of booking limits representing various ranges of constraint levels. For example, a 20% constraining level means that, on average, 20% of the data sets have their total demand constrained by the booking limit. To find the booking limits at these various levels, we use the *z*-score from a standard Normal distribution corresponding to the 20%, 40%, 60%, 80%, and 98% constrained levels, where *z* represents the number of standard deviations above or below zero for a standard Normal distribution. Thus, to find the *z*-score corresponding to 98%, we find the point where the area under the standard Normal curve equals 0.98, or z = 2.05. We then set our corresponding booking limits using:

$$BookingLimit = \mu + z * \sigma \qquad (6)$$

We test the five unconstraining methods across the three booking curve shapes (homogeneous, convex, and concave) for each of the five booking limits to test how each method performs under varied conditions. We chose unconstraining methods from previous research; the first three methods are the best performing methods from Weatherford and Polt's (2002) comparison. These include: 1) an averaging method (AM), called Naïve #3 by Weatherford and Polt, 2) Projection Detruncation (PD) and 3) Expectation Maximization (EM). The fourth method, lifetables (LT), is commonly used in medical studies and reliability engineering and is used in a revenue management context in van Ryzin and McGill (2000). A short description of each of these methods is provided in the appendix. The fifth method is DES, which was described in §3.

4.1 Results of Comparison

Overall, the EM and DES methods outperform the other three unconstraining methods. Table 2 shows the percentage error each unconstraining method produces compared to the actual mean of the demand distribution (the percentage errors were similar but slightly larger for the estimated variances). DES outperforms all methods for the homogeneous and convex data sets as its error remains less than 0.5% for all levels of constraining, compared to a maximum 5% error for the other methods. In the concave data set however, EM outperforms DES on average. Table 2 summarizes the results of the comparison and Figure 2 graphically summarizes the mean absolute error over all three curves.

Previous comparisons (Zeni, 2001; Weatherford and Polt, 2002) show EM outperforming PD; we confirm this result. Aside from the accuracy issues, PD has two disadvantages compared to EM: it takes more iterations to converge than EM and it requires the choice of a weighting parameter, τ, creating an opportunity for varying results. A $\tau < 0.5$ can lead to better results, but increases both the time to convergence and the chance for no convergence. As seen in Table 2, the averaging method is the worst performing, although in some instances the difference between it and the others is small. Thus, under some conditions, practitioners may be justified in forgoing the increased accuracy of the other methods for the computational simplicity of the averaging method.

Table 2 - Percentage error between calculated and actual mean for each unconstraining method

Booking Curve	Method	Percent of Data Sets Constrained				
		20%	40%	60%	80%	98%
Homogeneous	AM	-0.57%	-1.29%	-2.15%	-2.87%	-4.30%
	PD	0.23%	0.43%	0.56%	-0.53%	-2.99%
	EM	-0.06%	-0.23%	-0.56%	-0.22%	-0.58%
	LT	-0.17%	-1.31%	-1.53%	0.20%	0.43%
	DES	0.00%	0.00%	0.00%	0.00%	-0.14%
Convex	AM	-0.72%	-1.01%	-1.87%	-2.73%	-4.17%
	PD	0.26%	1.29%	0.98%	-0.91%	-2.99%
	EM	-0.08%	0.39%	-0.25%	-0.85%	-1.10%
	LT	0.13%	0.76%	0.89%	0.55%	5.72%
	DES	0.00%	0.14%	0.00%	-0.14%	-0.29%

Concave	AM	-0.71%	-1.43%	-1.86%	-2.57%	-4.57%
	PD	0.24%	0.73%	1.78%	-0.35%	-3.24%
	EM	-0.08%	-0.09%	0.09%	-0.19%	-0.93%
	LT	-0.21%	-0.83%	1.14%	-0.11%	2.29%
	DES	0.29%	0.71%	1.00%	1.86%	3.43%
Mean absolute error over all 3 Booking Curves	AM	0.67%	1.24%	1.96%	2.72%	4.35%
	PD	0.24%	0.82%	1.11%	0.60%	3.07%
	EM	0.07%	0.24%	0.30%	0.42%	0.87%
	LT	0.17%	0.97%	1.19%	0.29%	2.81%
	DES	0.10%	0.28%	0.33%	0.67%	1.29%

The lifetable method of unconstraining data produces estimates with errors very close to zero and even outperforms the DES and EM methods in a few of the concave cases. However, this method requires many computations and a large quantity of historical demand data. In a dynamic environment such as the travel industry, customer demand data changes quickly due to changes in the economic climate, broader market supply-demand-price relationships, and customer preference. Because of this, sufficient historical demand is often not available for the lifetable method to produce effective results.

For the homogeneous data pattern, the DES method has negligible error across the range of constrained data sets, due to the high predictability when arrival rates are constant over a given time period. For this data pattern, the DES method provides an estimate for the distribution mean that is up to 4% closer to the true mean than the next closes method. Similarly, for the convex data set, the DES method also provides the most accurate estimate in all the constraining conditions.

Figure 3 - Average absolute error from true demand for each unconstraining method

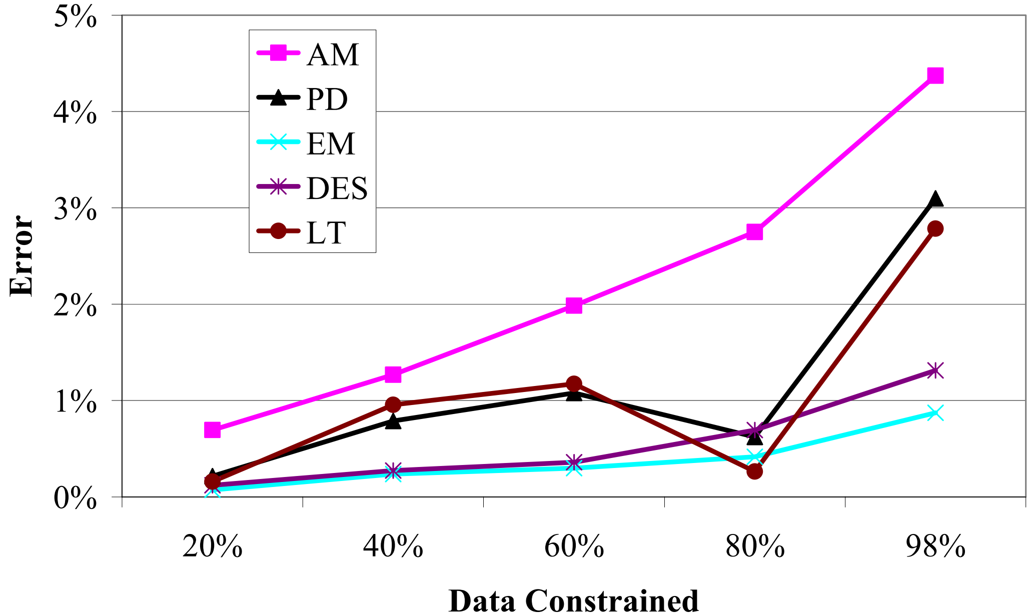

The DES method does not perform as well on the concave data set, although it still performs within a 1% error until demand is constrained in over 80% of the observations. The method underperforms on this demand pattern because booking segments close farther away from the arrival date for concave data, so many more data points must be estimated compared to the convex or homogeneous demand patterns. Here, the trend component of DES affects its accuracy as high demand occurring early in the booking curve is projected to continue once the booking limit has been met. For this booking curve shape, a forecasting method with a trend that is dampened over time may perform better. In practice however, an inaccuracy in unconstraining demand following a concave demand pattern is not a great concern. Airlines often offer cheaper fares to customers booking at least three weeks in advance. Because of this, and other similar restrictions, the lowest valued segment is often forced to follow the concave demand pattern. Due to the fundamental concepts behind revenue management, errors in estimating the true demand for the lowest valued segments are typically less costly than are errors in estimating the demand for higher valued segments.

4.2 Performance with Smaller Demand

The first set of results (Table 2) compares unconstraining methods when total demand averaged 698 units. However, in many applications, total demand is much smaller than 698, so we run a similar experiment with an average total demand of 19. We call this the "Small Demand" data set. Just as before, we ran simulations on homogeneous, concave, and convex booking curve shapes, with 100 trials of 40 days each for each shape. Results show that the DES and EM methods are the most accurate unconstraining method across a range of percentage constrained data, and the averaging method is the worst performing method. We notice the magnitude of error is much higher with the small demand data set as compared to the large demand data set due to difficulties of predicting data with intermittent demand (many periods with zero demand). This observation is consistent with previous studies; goods with intermittent demand are difficult to forecast and require specialized forecasting tools for the most effective results (Altay and Litteral, 2005).

Figure 4 - Unconstraining error with small demand data set

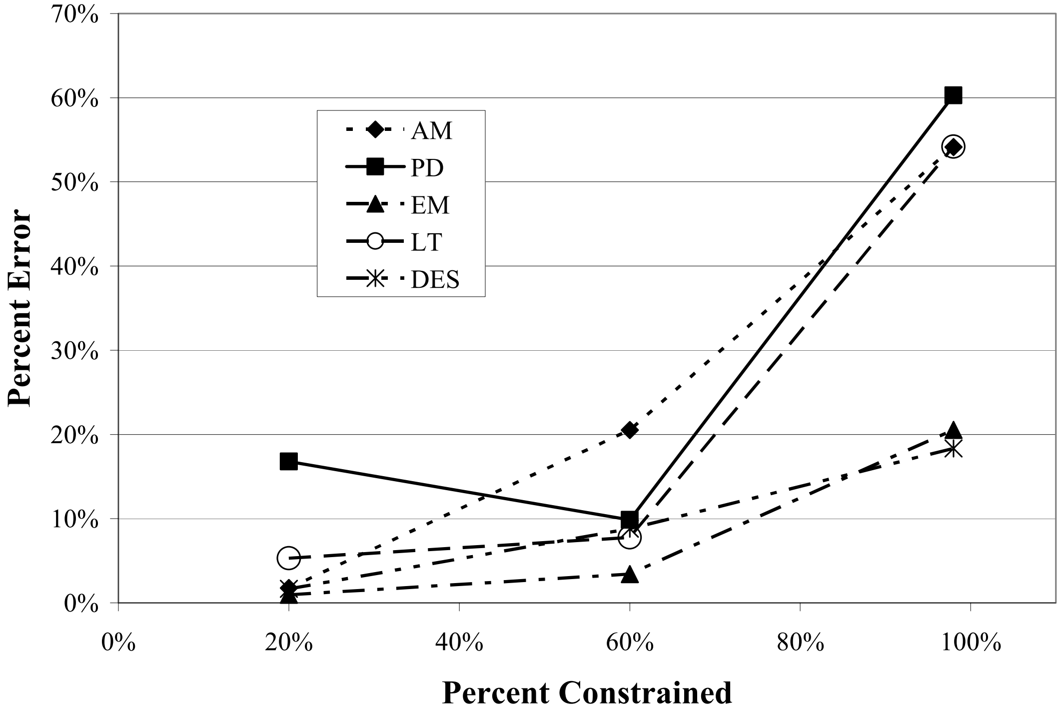

Figure 5 - Unconstraining error with large demand data set

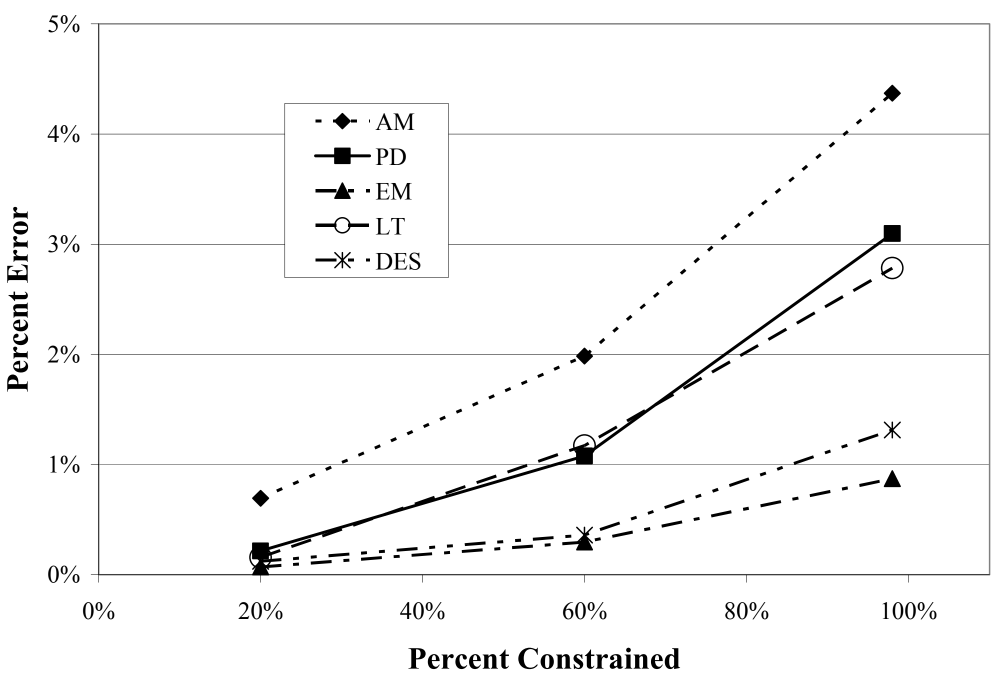

Since DES has significantly higher error with the small demand data set than with the large demand data set, an alternative formulation was sought. Croston's forecasting method (Croston, 1972) is a simple exponential smoothing method designed to accommodate small or intermittent demand. This method forecasts the size of the non-zero demands and inter-arrival time between non-zero demands. In a simple simulation over 60 trials and 60 days with a total demand of 12 (based on the smallest observed demand segment of our partner hotel), Croston's method outperforms DES across the range of constraints, as shown in Figure 6. These results provide evidence that Croston's method may be superior to DES for unconstraining when demand is intermittent. We have found that the conditions where Croston's method begins to outperform DES is when the percentage of days with zero demand exceeds 10% of the total number of days in the booking curve.

Figure 6 - Unconstraining error with small demand data set – DES vs. Croston's method

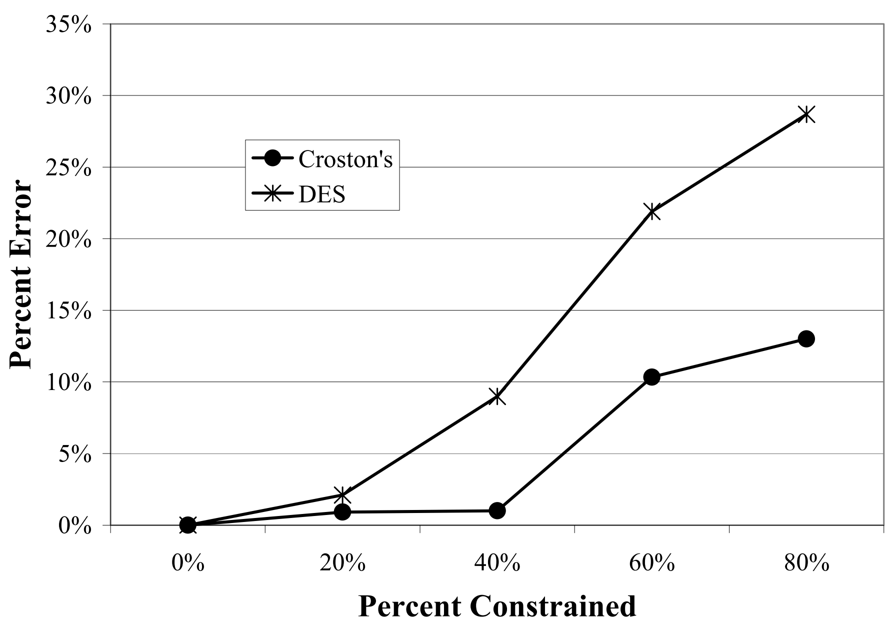

5. Revenue Impact Using Industry Data

In this section we compare the potential revenue impact of a major hotel/casino using DES versus EM versus no unconstraining. Since unconstraining methods only provide estimated parameters for the demand distribution, we use the EMSR-b (Expected Marginal Seat Revenue) algorithm (Belobaba, 1989) - a widely accepted method for setting booking limits for a basic revenue management system - to translate the demand distribution parameters (and corresponding room rates) into booking limits. The booking limits are then applied to booking data from a hotel/casino to calculate the total revenue impact. Thus, we compare the revenue convergence using EM, DES, and "ignoring unconstraining" based on actual (but normalized) booking and revenue data from a major hotel/casino. While the examples presented in this section are very useful for illustrating the effectiveness of the methods, they cannot lead to conclusions about industry performance. Such conclusions can only be drawn from trials in practice.

5.1 Demand Data

We use actual hotel/casino booking data to test the impact that unconstraining has on revenue. We use booking curve data for 12 consecutive Friday night stays, unconstrained using direct observation of turndowns. Extra care was taken to ensure that all demand was captured for this data set including demand that occurred after booking limits were met. Because of the increased cost involved in such careful data collection, we limited the data collection period to 12 weeks and used bootstrapping to create 1000 booking curve samples from the initial data. Hotel reservations vary greatly by day of the week, depending on the type of hotel. For this hotel, weekends are the most popular, and therefore, have the highest constraining rate. In order to control for differences in demand between different days of the week, we focus on Friday night stays during the 12-week period. Within any Friday night's booking data, this hotel/casino has many different customer segments, with some customer segments so valuable they are rarely constrained (revenue per night from the highest fare customer can be 12 times the revenue from the lowest fare customer); therefore we focused our unconstraining efforts on the most popular four segments that are constrained.

Bootstrapping was performed as follows. First, observing the 12 Friday night stay booking curves, it was apparent that at different intervals before arrival, the slope of the curve changes dramatically. Based on these slope changes, we created multiple intervals within the 60 day window which had similar arrival rates. Picking randomly (with replacement) from 12 weeks' worth of Friday night booking patterns within a similar arrival rate interval, we use the bootstrap method to create 1000 different 60 day booking curves for each of the four customer segments.

Observing the original hotel booking data showed that the hotel would temporarily close the lower valued segments midway through the 60 day booking curve, thus we simulated this practice of opening and closing the classes multiple times. We closed a booking class (constrained demand) midway through the booking curve, reopened the booking class, and then closed it again before the actual day of arrival. We did this for each of the 4 different segments in all 1000 replications. We set protection levels so that 50% of a given data set would be constrained, then again so that 75% of a given data set would be constrained, to test our methodology against different constraining levels. Using both DES and EM, we unconstrained these data sets and compared the distribution parameter estimates for each method against the true parameter values. Both methods performed well, with average errors listed in Table 3. Just as in previous trials, all methods perform better with less constrained data. The methods better predict mean values than standard deviations. Over 1000 instances, both methods predict the mean within 5% of the true mean, showing the methods perform well even when demand is constrained multiple times in a booking curve.

Table 3 – Error comparison between EM and DES with interrupted arrivals and actual data

		50% of Data Sets Constrained	75% of Data Sets Constrained
Mean	EM	0.84%	1.72%
	DES	0.84%	4.87%
Standard Deviation	EM	7.38%	14.72%
	DES	7.65%	23.86%

Next we testthe impact each unconstraining method would have on the hotel's revenue. Total revenue from a revenue management system is the ultimate indicator of a system's success. Unconstraining methods however, only provide estimates for the demand distribution parameters. Thus, we borrow van Ryzin and McGill's (2000) general methodology for translating protection levels into revenue. To test convergence and robustness, we start with purposefully high and low protection levels, similarly to van Ryzin and McGill (2000). Protection levels must be set at some estimated level for initial product offerings because a firm will have very little idea of demand for a new product (protection levels are the opposite of booking limits, i.e. how many units of capacity at a given class to protect for higher fare classes). As more demand is observed over time, the firm adjusts protection levels accordingly to increase total revenue. The convergence rate to optimal protection levels depends on both the starting levels chosen and the unconstraining method used. Thus, we test protection level convergence and total revenue convergence using the two best performing unconstraining methods (EM and DES) with two different starting protection levels – low and high. To underscore the importance of unconstraining, we included data with no unconstraining, labeled (Spiral) for the "spiral-down effect" which occurs when data is not unconstrained.

5.2 Setting Protection Levels: the EMSR-b Method

For setting the protection levels for the hotel rooms, we use a variation of the EMSR (Expected Marginal Seat Revenue) heuristic (Belobaba 1989), called EMSR-b. This is the most common heuristic used in practice for setting protection levels. The EMSR-b method does not produce optimal protection levels under all real world conditions, but is representative of a basic revenue management system and is sufficient for comparing unconstraining methods. EMSR-b works as follows: Given the estimates of the means, $\hat{\mu}_i$ and standard deviations, $\hat{\sigma}_i$ for each customer value segment *i*, the EMSR-b heuristic sets protection level θ_i so that $f_{i+1} = \bar{f}_i P(\overline{X}_i > \theta_i)$, where $\overline{X}_i$ is a normal random variable with mean $\sum_{j=1}^{i} \hat{\mu}_j$ and variance $\sum_{j=1}^{i} \hat{\sigma}_j^2$, f_i is the revenue for customer value segment *i* and $\bar{f}_i$ is weighted average revenue, given by:

$$\bar{f}_i = \frac{\sum_{j=1}^{i} f_j \hat{\mu}_j}{\sum_{j=1}^{i} \hat{\mu}_j} \tag{7}$$

In simpler terms, this rule performs a marginal analysis on the benefits of holding capacity for a higher valued customer versus the cost of turning away the next lower valued customer. However, protection levels can only be optimized if the true demand distribution parameters are known; hence the need for a good unconstraining method.

5.3 Simulation

We test the revenue impact of the unconstraining methods by applying protection levels (based on the EMSR-b method using the distribution parameter estimates from the unconstraining method) to the industry data described in §5.1. Total revenue is calculated by multiplying the customer value of each segment by the number of reservations that would have been sold in that segment; the min of the protection level and the total demand. Mathematically, this is shown as:

$$\text{Revenue} = \sum_{i=1}^{4} V_i * Min\left[\theta_i, x_i\right] \qquad (8)$$

where

V_i = customer value for segment i

x_i = total demand for segment i

θ_i = booking level for segment i

Normalizing the fare class data from our hotel/casino on a scale from $1 - $100, the per-night expected revenues for the four customer segments are: $25, $35, $62, $100. These expected revenues are based on the total amount a customer in that segment is expected to spend at the hotel/casino per night, including the revenue from the room rate, food, beverages, shows, and casino. Customer spending is tracked over time by issuing frequent stay cards that are recorded each time the customer makes a transaction.

Since estimates for the parameter values of the demand distribution, and the corresponding protection levels, evolve over time, we simulate this evolution in our study. First, we split each of the 4 sets of 1000 booking curves described in §5.1 into 10 sets of 100 booking curves Working with the first set of 100 booking curves, we estimated initial protection levels for each customer segment (two initial starting protection levels are used for each segment, one lower and the other higher than the optimal protection levels). We then calculated the revenue the hotel/casino would have received if they used these initial protection levels for each segment over all 100 booking curves in the set. The first data point in Figures 7 and 8 is the percentage difference in revenues the hotel/casino would have received using these initial protection levels versus if they had used optimal protection levels calculated with the true demand distribution parameters. Next, we applied the initial protection levels to the first

10 booking curves (booking curves 1-10 out of the 100 in the set). Thus, some segments of the first 10 booking curves were constrained by the calculated protection levels. We applied each of the unconstraining methods to this group of 10 constrained booking curves and calculated new protection levels for the next group of 10 booking curves (booking curves 11-20 out of the 100 in the set). Based on these new protection levels, we calculated the revenue generated if these protection levels were used on all 100 booking curves in the set. This procedure continued, unconstraining the demand data and readjusting the protection levels every 10 booking curves, until all 100 booking curves in the set were used. This procedure simulates a hotel manager watching demand for 10 consecutive Fridays, then adjusting his protection levels for the next 10 Fridays, and continuing this procedure for a total of 100 consecutive Fridays.

For robustness, we applied this methodology to 10 sets of 100 booking curves in order to calculate a standard error of our estimates. The data plotted in Figures 7 and 8 represent an average over the 10 sets along with upper and lower limits corresponding to a 95% confidence interval. All of the methods (EM, DES, Spiral) were tested against the same data sets (the simulations are coupled)

Revenues for each of the unconstraining methods are compared to optimal revenue. We calculate the optimal revenue by finding the mean and standard deviation of each set of 1000 booking curves (one set of 1000 for each of four customer segments). Since the initial data was unconstrained, we knew the true demand for every booking curve, and hence knew the true mean and standard deviation parameter estimates. We applied EMSR-b using these parameters to find protection levels, then using nested protection levels, applied (8) to calculate total revenue.

Figure 7 – Revenue achieved using high protection limits for EM, DES, and no unconstraining

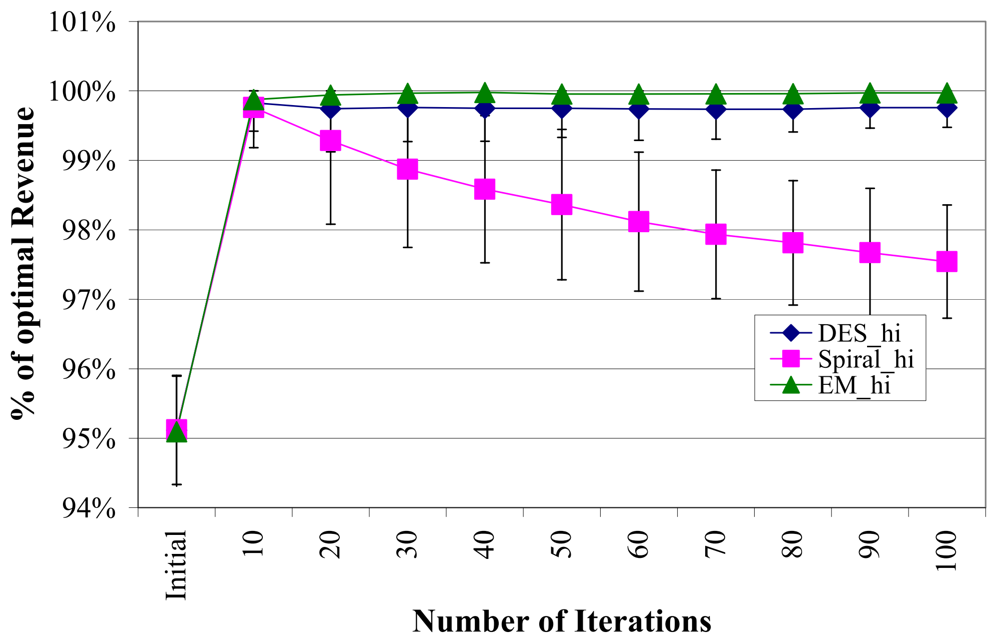

Figure 7 shows a convergence to the optimal revenue for DES and EM methods after starting with the high initial protection levels. Here both unconstraining methods allow the EMSR-b method to quickly converge to optimal protection levels and hence achieve optimal revenues. The DES and EM methods yield similar results, both starting at 95.5% of optimal revenues and improving to close to100% after only one iteration. High starting protection levels restrict early bookings in the lower value segments while saving capacity for the high value segments. When historical data is limited and the difference in revenue between high and low value segments is large, a firm may want to initially employ high protection levels.

Figure 8 - Revenue achieved using low protection limits for EM, DES, and no unconstraining

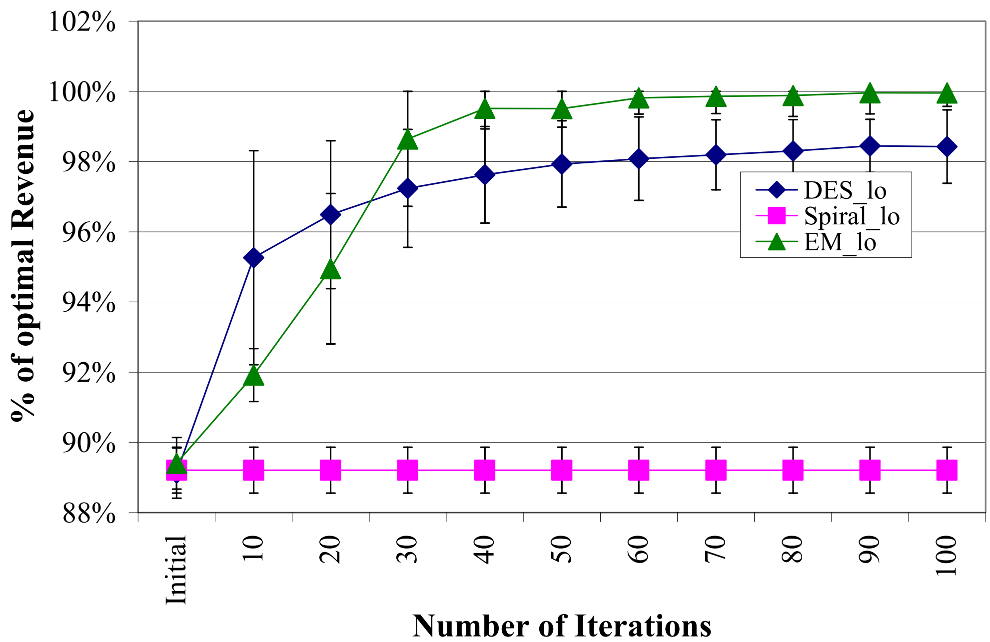

Compared to the high starting protection levels in Figure 7, the low starting protection levels in 8 converge to the optimal revenue for both unconstraining methods at a much slower rate, as all three of the highest fare classes are initially 100% constrained. Because of this, we were unable to use the EM method during the first iteration since the EM method requires at least one unconstrained booking curve. Instead, we increased the protection limits by 10% for each group of 10 booking curves until at least one booking curve was unconstrained, at which point we could begin using the EM method. While this practice may seem arbitrary (justifiably so), it is representative of techniques commonly applied in practice. The DES method does not suffer from such a limitation, thus it outperforms the EM method during the early stages of the low starting protection levels case. Once the EM method can be used, revenue quickly converges to greater than 99% of optimal revenue. DES performs better in early iterations, then converges to just above 98% of optimal revenues, slightly trailing the EM method's performance.

Both Figures 7 and 8 illustrate that unconstrained data may lead to a loss in revenue. When starting with low protection levels, Figure 8 shows the "no unconstraining" (Spiral) data never improves past the 89% of revenue initially achieved. This compares unfavorably with the two unconstraining methods which steadily improve as more demand is observed. Figure 7 shows that when using high initial protection levels, failing to unconstrain data causes revenue to decrease every time protection limits are recalculated as the historical data becomes more and more constrained. This shows graphically the "spiral-down effect" described by Cooper et al. (2005).

One may conjecture from a comparison of Figures 7 and 8 that, in the absence of historical demand data, it is always better to start with high protection levels versus low since the revenue converges to the optimal much faster in Figure 7 than in Figure 8. Such a generalization is incorrect however, as this phenomenon is an artifact of our choice of revenues for each customer class. For this hotel, the difference in revenues between the highest fare class ($100) and the next highest ($62) is much larger than the difference between the two lowest fare classes ($35 and $25). If these differences had been reversed (say fare classes of $100, $90, $63, and $25), starting with low initial protection levels would converge to the optimal revenue much faster than starting with high protection levels.

We note a few additional observations from our bootstrapping results. First, when demand has sudden shifts or a large number of constrained days, all statistical methods become less accurate. In practice, we recommend qualitative adjustments to the demand data or the numerical forecast for these situations. Second, in any situation, statistical unconstraining should be supplemented by a physical constrained count, where possible, to check the validity of unconstrained forecasts. For a hotel, this physical count could include reservation agents and bellhops keeping a manual tally of the number of people turned away. A store might better promote "rain checks" of sold out items and keep track of how many people ask for the sold out good. Third, EM becomes more accurate for larger sample sizes, but often performs poorly for small sample sizes. In these cases, DES may be a better solution until more historical data is available. Fourth, a small firm with a limited IT budget may not have the resources to afford a sophisticated statistical program to run the EM method. Such a firm may find it more useful to use the straightforward DES method to estimate total demand. Lastly, if data is fully constrained, EM does not work, and an alternate method must be used. Either an alternative unconstraining method or a "rule of thumb" adjustment to the protection level is needed until unconstrained demand is observed.

6. Conclusions/ Recommendations

This paper examines the often overlooked but essential topic of unconstraining sales data. True demand distribution parameters are a critical ingredient to revenue management systems; unfortunately the data available is often constrained. Ignoring the constrained data problem results in significant reductions in revenue and observing demand after it exceeds capacity is often impractical, thus statistical unconstraining methods are often used to estimate the parameters of the demand distribution. We propose a new unconstraining method (DES) based on a common forecasting model that, unlike traditional statistical unconstraining methods, takes into account the point in time on the historical demand booking curves that demand was constrained. We find that our proposed DES method and the EM method perform similarly well and outperform the alternatives. When little historical data is available or all demand sets are constrained, then DES is an better choice than EM. Also, if sophisticated statistical software packages are not available, DES provides a better alternative than the averaging method.

We test the revenue impact of DES, EM, and "no unconstraining" on actual booking data from a hotel/casino. We show that performance varies with the initial starting protection limits and a lack of unconstraining leads to significant revenue losses. In our example, starting with high(low) initial protection levels converge to optimal limits more quickly(slowly). Low initial protection levels can lead to completely constrained classes, forcing firms using the EM method for unconstraining to use other methods until some unconstrained booking curves are observed. Both EM and DES take numerous iterations (5 and 3, respectively) to converge to within 2% of optimal with low initial protection levels. On the other hand, both EM and DES converge to within 0.3% of optimal after only one iteration when starting with high protection levels. These results show that both EM and DES are effective methods for unconstraining and provide much higher revenues than no unconstraining at all.

As is true with all research, there are limitations to our work. In our section on revenue impact, we assume independent demand for a given customer value segment. That is, a customer associated with value segment 2 will not turn up as demand in segments 1 or 3, even if value segment 2 is closed. This is a very realistic assumption in a casino application, where customer gaming habits are not dependent upon the rate program they book under, but rather vice versa. However, in many other revenue management applications – particularly airlines - demand for a given segment often depends on the choice of segments available. The independence assumption is commonly used however in most booking limit algorithms. To incorporate consumer choice behavior such as buy-up or buy-down behavior, Talluri and van Ryzin (2004a) present a model that explicitly accounts for the probabilities that customers in a given fare class (value segment) will purchase from other fare classes if their preferred fare class is unavailable. To use this model, however, a firm needs to know the probabilities that customers in all classes will buy-up, or buy-down, probabilities that are rarely known in practice. Further development of such models along with unconstraining methods to accommodate them is a promising area of future research. Additionally, our DES method does not account for seasonality or price promotions. Like most other forecasting methods, historical data should be decomposed into components of promotion effects, seasonality, and competitive effects before DES is applied. For further information on these adjustments, see Armstrong (2001).

References

Altay, N. and Litteral, L. 2005. "A Comparative Study of Intermittent Demand Forecasting Models", working paper, Robins School of Business, University of Richmond.

Armstrong, J. Scott (ed.). 2001. *Principles of Forecasting: A Handbook for Reseracher and Practitioners*. Norwell, MA: Kluwer Academic Publishers.

Badinelli, R.D. 2000. "An Optimal, Dynamic Policy for Hotel Yield Management," *European Journal of Operational Research* **121**: 476-503.

Belobaba, P.P. 1989. "Application of a Probabilistic Decision Model to Airline Seat Inventory Control," *Operations Research* **37** (2): 183-197.

Bitran, G.R. and S. M. Gilbert. 1996. "Managing Hotel Reservations with Uncertain Arrivals," *Operations Research* **44**: 35-49.

Bitran, G.R. and S. V. Mondschein. 1995. "An Application of Yield Management to the Hotel Industry Consdiering Multiple Day Stays," *Operations Research* **43**: 427-443.

Carroll, B. and J. Siguaw. 2003. "The Evolution of Electronic Distribution: Effects on Hotels and Intermediaries," *Cornell Hotel and Restaurant Administration Quarterly* **44**(4) 38-50.

Cooper, W.L, T. Homem-de-Mello, A.J. Kleywegt. 2005. "Models of the Spiral-Down Effect in Revenue Management," Working Paper, Department of Mechanical Engineering, University of Minnesota.

Cox, D.R. 1972. "Regression Models and Life Tables (with discussion)," *Journal of the Royal Statistical Society*. **B34** (2): 187-220.

Cox, D.R. and D. Oakes. 1984. *Analysis of Survival Data*. New York, Chapman & Hall.

Cross, R. G. 1997. *Revenue Management*. New York, Broadway books.

Croston, J.D. 1972. "Forecasting and Stock Control for Intermittent Demands," *Operational Research Quarterly* **23** (3): 289-303.

Dempster, A.P., N.M. Laird, and D.B. Rubin. 1977. "Maximum Likelihood from Incomplete Data via the EM Algorithm," *Journal of the Royal Statistical Society*, *Series B*. **39** (1): 1-38.

Kalbfleisch, J.D. and R.L. Prentice. 1980. *The Statistical Analysis of Failure Time Data*. New York, John Wiley.

Lawless, J.F. 1982. *Statistical Models and Methods for Lifetime Data*. New York, John Wiley.

Liu, P.H. and V. Makis. 1996. "Cutting-tool reliability assessment in variable machining conditions," *IEEE Transactions on Reliability*. **45**(4): 573-581.

Liu, P.H, S. Smith, E.B. Orkin, and G. Carey. 2002. "Estimating unconstrained hotel demand based on censored booking data," *Journal of Revenue and pricing Management*. **1**(2): 121-138.

Maddala, G.S. 1983. *Limited Dependent And Qualitative Variables in Econometrics*. Cambridge, Cambridge University Press.

McGill, J.I. 1995. "Censored regressed analysis of multiclass passenger demand data subject to joint capacity constraints," *Annals of Operations Research* **60**:209-240.

McGill, J.I. and G.J. van Ryzin. 1999. "Revenue Management: Research Overview and Prospects," *Transportation Science* **33**(2): 233-256.

Nelson, W. 1990. *Accelerated Testing: Statistical Models, Test Plans, and Data Analyses*. New York, John Wiley.

Orkin, E.B. 1998. "Wishful Thinking and Rocket Science," *Cornell Hotel and Restaurant Administration Quarterly* **39**(4) 15-20.

Rothstein, M. 1974. "Hotel Overbooking as a Markovian Sequential Decision Process," *Decision Sciences* **5**:389-404.

Schneider,H . 1986. *Trucated and Censored Samples from Normal Populations*. New York, Marcel Dekker.

Talluri, K. and G.J. van Ryzin. 2004a. "Revenue Management Under a General Discrete Choice Model of Consumer Behavior," *Management Science* **50** (1): 15-33.

Talluri, K. and G.J. van Ryzin. 2004b. The Theory and Practice of Revenue Management, Kluwer Academic Publishers, Norwell, Massachusetts.

van Ryzin, G.J. and J.I. McGill. 2000. "Revenue Management Without Forecasting or Optimization: An Adaptive Algorithm for Determining Airline Seat Protection Levels," *Management Science* **50** (1): 15-33.

Weatherford, L.R. and S.E. Bodily. 1992. "A Taxonomy and Research Overview of Perishable-Asset Revenue Management: Yield Management, Overbooking, and Pricing," *Operations Research* **40**(5):831-844.

Weatherford, L.R. and S. Polt. 2002. "Better Unconstraining of Airline Demand Data in Revenue Management Systems for Improved Forecast Accuracy and Greater Revenues," *Journal of Revenue and Pricing Management* **1**(3):234-254.

Zeni, R. 2001. "Improving Forecast Accuracy by Unconstraining Censored Demand Data" AGIFORS, Bangkok, Thailand.

Appendix: Description of Unconstraining Methods

Averaging Method (AM) – As described by Weatherford and Polt (2002), "replace 'closed' observations with the larger of (a) actual observation or (b) average of the 'open' observations. Calculate average and standard deviation based on open observations and modified values of closed observations.

Expectation Maximization (EM) – As explained in Talluri and van Ryzin (2004b) –
Suppose we have $M+N$ observations of bookings for a given product, $z_1, \ldots z_{M+N}$, of which M observations are constrained because the product was closed. We ignore the time-series aspect of the observations and treat $z_1, \ldots z_{M+N}$ as an unordered set of observations generated by an i.i.d. process….. Our goal is to find the parameters of an underlying demand distribution for these observations.

Assume that the underlying demand distribution is normal with mean μ and standard deviation σ. We further assume that all the observations come from a common distribution and that the observations are constrained at random, i.e., they appear randomly in the sample. Since we are treating the observations as unordered, assume $z_1,\ldots.z_M$ are constrained at booking limits $b_1,\ldots.,b_M$, so that $z_1=b_1,\ldots,z_M=b_M$. The remaining N observations are unconstrained.

If the data were not constrained, then it would be easy to construct the complete-data likelihood function. Namely,

$$L(\mu,\sigma,M+N)=\prod_{i=1}^{M+N}\frac{1}{\sqrt{2\pi\sigma}}e^{\frac{-(z_i-\mu)^2}{2\sigma^2}}$$

with the complete-data log-likelihood function given by

$$\ln L(\mu,\sigma,M+N)=-\frac{M+N}{2}\ln 2\pi-(M+N)\ln\sigma-\frac{\sum_{i=1}^{M=N}(z_i-\mu)^2}{2\sigma^2}$$

The μ and σ that maximizes ln $L($ $)$ in the above equation are given by the closed-form solution

$$\hat{\mu}=\frac{1}{M+N}\sum_{i=1}^{M+N}z_i$$

$$\hat{\sigma}^2=\frac{1}{M+N}\sum_{i=1}^{M+N}(z_i-\hat{\mu})^2$$

However, we do not know the true values of the M constrained observations $z_1,\ldots.z_M$ and therefore cannot use this procedure directly.
We use the following steps:

Step 0 (Initialize): Initialize μ and σ to be $\mu^{(0)}$ and $\sigma^{(0)}$. Good candidates for these starting values are the sample mean and sample standard deviation of all the unconstrained observations. Let $\delta>0$ be a small number, to be used as a stopping criterion.

$$\mu^{(0)} = \frac{\sum_{i=M+1}^{M+N} z_i}{N}$$

$$\sigma^{(0)} = \sqrt{\frac{\sum_{i=M+1}^{M=N} (z_i - \mu^{(0)})^2}{N}}$$

Step 1 (E-step): Calculate the expected value of the censored data in the log-likelihood function assuming that they come from a normal distribution X with parameters ($\mu^{(k-1)}$, $\sigma^{(k-1)}$). That is, for i = 1,….,M calculate

$\hat{Z}_i^{(k-1)} \doteq E[X \mid X \geq b_i, X \sim N(\mu^{(k-1)}, \sigma^{(k-1)})]$ and

$$(\hat{Z}_i^2)^{(k-1)} \doteq E[X^2 \mid X \geq b_i, X \sim N(\mu^{(k-1)}, \sigma^{(k-1)})]$$

Next, for each censored observation i = 1, …, M, replace z_i by $\hat{Z}_i^{(k-1)}$ and z_i^2 by $(\hat{Z}_i^2)^{(k-1)}$ to form the complete-data log-likelihood function Q(μ,σ). In this way, we replace the constrained values in the log-likelihood function by their expected alues given the current estimates of the mean and standard deviation.

Step 2 (M-step): Maximize Q(μ,σ) with respect to μ and σ to obtain $\mu^{(k)}$, $\sigma^{(k)}$, yielding

$\mu^{(k)} = \frac{1}{M+N}\left[\sum_{i=1}^{M} \hat{Z}_i^{(k-1)} + \sum_{i=M+1}^{M+N} z_i\right]$ and

$$\sigma^{(k)} = \frac{1}{M+N}\left[\sum_{i=M+1}^{M+N}\left(\left(\hat{Z}_i^2\right)^{(k-1)} - 2\hat{Z}_i\mu^{(k-1)} + \left(\mu^{(k-1)}\right)^2\right) + \sum_{i=M+1}^{M+N}\left(z_i - \mu^{(k-1)}\right)^2\right]$$

Step 3 (Convergence test): If $\| \mu^{(k)} - \mu^{(k-1)} \| < \delta$ and $\| \sigma^{(k)} - \sigma^{(k-1)} \| < \delta$, then stop; else, $k \leftarrow k+1$, goto step 1.

Once convergence has been achieved – say, in iteration K – the unconstrained values for z_i, i = 1, …, M can be taken as $E[X \mid X \geq b_i]$, where X is normally distributed with mean $\mu^{(K)}$ and standard deviation $\sigma^{(K)}$.

Projection Detruncation (PD) – as explained in Weatherford and Polt (2002) – This is a variation of the EM method, which uses a weighted value, τ, to scale the aggressiveness of unconstraining. Instead of calculating a conditional expectation, it iteratively finds a mean and standard deviation such that

τ = area under normal curve from unconstrained mean to infinity / area under normal curve from constrained mean to infinity

The user chooses τ, where $0 \leq \tau \leq 1$. τ=0.5 performs similarly to EM, using the conditional median rather than a conditional mean. $\tau \leq 0.5$ produces larger estimates for distribution parameters than $\tau \geq 0.5$.

Lifetables (LT) – as described by Talluri and van Ryzin (2000)

Let n denote the toal number of observations (censored and uncensored). Let $t_1 < t_2 < \ldots, t_m$ be m distinct intervals. (Call $[t_j, t_{j+1})$ interval j.) Let n_j be the number of observations with values t_j or more (the number of "at risk" obersvations at the start of interval j); let d_j be the number of uncensored observations that fall in interval j (the number of "deaths" in the interval j); and let w_j be the number of censored observations that fall in interval j (the number of "withdrawals" because of censoring in interval j). Define $n_0 = n$ and note that $n_j = n_{j-1} - d_j - w_j, j = 1, \ldots, m$.

The life table estimate is given by

$$\hat{S}\left(\frac{t_j + t_{j+1}}{2}\right) = \prod_{i=1}^{j}\left(1 - \frac{d_i}{n_i - w_i / 2}\right)$$

Each term $1 - (d_j/(n_i - w_i/2)$ is an estimate of the conditional probability that demand exceeds t_{i+1} given that it exceeded t_i. The denominator, $n_i - w_i/2$, is an estimate of the number of samples at risk during period i, where $w_i/2$ is a correction term for the number of censored observations in period i. We used 20 intervals in our work.

We then used the life table estimator to estimate the mean and standard deviation of the distribution by linear regression. Specifically, let $\Phi(x)$ be the standard normal distribution and let $\Phi^{-1}(x)$ denote its inverse. Define $s_j = \Phi^{-1}(1 - \hat{S}((t_j = t_{j+1})/2))$. If demand is normally distributed, the points (s_j, t_j) j = 1, …, n should lie approximately on a straight line, namely $s_j = at_j + b$. Using linear regression, we use our estimates of the slope, $\hat{a}$, and the intercept, $\hat{b}$, to estimate the mean, $\hat{\mu} = 1/\hat{a}$, and the standard deviation, $\hat{\sigma} = -\hat{b}/\hat{a}$.